

Topics in Interoperability: Infrastructure Replacement in a System of Systems

David Carney
James Smith
Patrick Place

November 2005

Integration of Software-Intensive Systems Initiative

Unlimited distribution subject to the copyright.

Technical Note
CMU/SEI-2005-TN-031

This work is sponsored by the U.S. Department of Defense.

The Software Engineering Institute is a federally funded research and development center sponsored by the U.S. Department of Defense.

Copyright 2005 Carnegie Mellon University.

NO WARRANTY

THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

Use of any trademarks in this report is not intended in any way to infringe on the rights of the trademark holder.

Internal use. Permission to reproduce this document and to prepare derivative works from this document for internal use is granted, provided the copyright and "No Warranty" statements are included with all reproductions and derivative works.

External use. Requests for permission to reproduce this document or prepare derivative works of this document for external and commercial use should be addressed to the SEI Licensing Agent.

This work was created in the performance of Federal Government Contract Number FA8721-00-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. The Government of the United States has a royalty-free government-purpose license to use, duplicate, or disclose the work, in whole or in part and in any manner, and to have or permit others to do so, for government purposes pursuant to the copyright license under the clause at 252.227-7013.

For information about purchasing paper copies of SEI reports, please visit the publications portion of our Web site (<http://www.sei.cmu.edu/publications/pubweb.html>).

Contents

Abstract.....	v
1 Introduction.....	1
1.1 Interoperation as a Relationship.....	1
1.2 Boundaries of Systems and Systems of Systems.....	1
1.3 Relationships Implemented by Systems.....	3
2 Context and Background	4
2.1 Description of the COS.....	4
2.2 Organizations and Their Missions	6
2.3 Interoperability in the COS.....	6
2.4 Description of the Planned Upgrade.....	7
3 Specific Interoperability Issues.....	9
3.1 System Interfaces.....	9
3.1.1 Elements of Risk.....	9
3.1.2 Observed Mitigations	10
3.2 Organizational Responsibilities.....	11
3.2.1 Elements of Risk.....	12
3.2.2 Observed Mitigations	12
3.3 Requirements and Functionality	13
3.3.1 Elements of Risk.....	14
3.3.2 Observed Mitigations	14
3.4 Development and Integration Processes.....	14
3.4.1 Elements of Risk.....	14
3.4.2 Observed Mitigations	15
3.5 Testing	15
3.5.1 Elements of Risk.....	15
3.5.2 Observed Mitigations	16
4 Summary	17
References.....	19

List of Figures

Figure 1: Systems and Systems of Systems	2
Figure 2: High-Level COS Organization	5
Figure 3: Organizational Links	11

Abstract

This technical note examines the Common Operations System (COS), a large aggregation of independently developed systems, and the risks posed to it by an infrastructure upgrade. Many large organizations involved in various critical government roles depend on the COS for planning their business operations. When such a large number of applications rely on a complex infrastructure, an attempt to upgrade raises many interoperability issues. The risks involved, and their observed mitigations, are examined in several areas: system interfaces, organizational responsibilities, requirements and functionality, developing an integration process, and testing.

1 Introduction

This technical note is a case study on replacement of the infrastructure of a large system of systems. We begin with a brief, general discussion of some key concepts that clarify our understanding of interoperability. While this discussion is not directly relevant to the remainder of the report, we believe that this conceptual introduction will be useful to the reader's understanding of our analyses of the system described in this report.

1.1 Interoperation as a Relationship

The term *interoperability* has many definitions; a reasonable one is

the ability of a collection of communicating entities to (a) share specified information and (b) operate on that information according to a shared operational semantics in order to achieve a specified purpose in a given context [Carney 05].

The essence of interoperation is that it is a **relationship** between **systems**, where systems are the entities in the above definition. While our focus will be on computer-based systems, the definition extends beyond the world of mechanical systems to organizational and other contexts. To interoperate, one system must provide a service¹ that is used by another. This cannot be achieved without, at a minimum, communication, whether direct or indirect, from the provider to the consumer of the service.

Two or more systems that have interoperability relationships form a **system of systems**. We suggest that the following characteristics, adopted from Maier, will be exhibited by any system of systems [Maier 98]:

- operational **and** managerial independence of the elements
- evolutionary development
- emergent behavior

1.2 Boundaries of Systems and Systems of Systems

Almost every discussion of interoperability is plagued by one annoying reality: any construct that we label a *system* may in fact be composed of several constituent systems, and this may recursively be true at several levels. In other words, anything that at one level we can call a

¹ While it seems obvious, it must be stated that provision of service includes the provision of data.

“system” may internally be a “system of systems,” and any “system of systems” may itself be part of some larger “system of {systems of systems},” and so forth.

To illustrate, we imagine some hypothetical data systems that interoperate in some manner. These data systems could all be elements (e.g., communication or navigation) of a military aircraft’s avionics system, which together with many other systems (weapons system, mission management system) compose the total aircraft, which itself can be viewed as a single system. To continue to even higher levels, the aircraft is an element in a larger system of systems, since it interoperates with other aircraft and other military units in combat. The process can continue recursively through ever larger systems of systems of systems of systems.

To facilitate our discussion of interoperability, we need to define some level of immediate interest. To do so we choose one of these many levels as that of “the system” and the next higher level as that of the “system of systems.” The level we choose is, in a sense, arbitrary, since it is only one vantage point within the potentially large scope of this recursive sequence. But it is useful to focus discussion and analysis.

Thus, if our concern at the moment is with issues related to low-level data, semantics of data, and so forth, we could choose the data systems noted above and their interoperability relationships as our level of interest.

We illustrate this as follows:




Figure 1: Systems and Systems of Systems

Graphically, therefore, we let the three smaller circles above represent the individual data systems in the hypothetical example described earlier. Each is related to two others by some

interoperability relationship. The three together as a related unit, that is, the “system of systems,” is depicted by the large darker oval; this would be the hypothetical avionics system. The smaller circles may themselves each comprise several systems, and the large oval may itself be a single system in some larger context. We temporarily ignore those possibilities and focus only on the interrelationships between A, B, and C that bring about D. By choosing this particular vantage point, we are able to consider the precise nature of the three constituent systems, their interrelationships, and the principles by which the system of systems (D) is brought about.

At some later time, if our concern lies in some other sphere (e.g., real-time factors relating to the avionics and weapons systems), then our level of discourse could well be the interoperability relationships at that level, and so on.

1.3 Relationships Implemented by Systems

A further facilitating device is that we use a common vocabulary regardless of the mechanism by which a relationship is implemented. For example, we can imagine two systems (A and B) whose relationship is such that they must communicate data back and forth. Let us further suppose that the relationship is implemented by some complex communication system. Since that communication system is, by definition, a system in its own right, it is easy to see that discussion of such a collection may easily be complicated by two different opinions. One opinion sees a system of systems of three entities (A, B, and the communication system). The other opinion sees a system of systems of only two (i.e., by disregarding that the communication system is a **system**, and viewing it only as implementing the relationship between A and B).

We argue that either view is possible, depending on the issues of immediate interest and what questions are being asked. For instance, we may be interested in the semantics of shared data between A and B, and are unconcerned with the manner in which the data is communicated. In that case, we can rightly consider the communication system simply as the mechanism that implements the A-B relationship. On the other hand, if we are concerned with the specific details of how System A locates System B, with the significance of timing constraints and other such questions, then we well may consider that the relationships between System A, System B, and the communication system are all interoperability relationships in their own right.

We now turn to examination of a large and complex system of systems that has been in operation for several years. The study focuses on the planned replacement of its key infrastructure system, and particularly on the effect this will likely have on the interoperability relationships that now exist throughout the overall system of systems. The proposed upgrade is planned for implementation within the next two years. Section 2 describes the context and background of the system(s) and the organizations involved in the planned upgrade. Section 3 examines several specific interoperability issues that are critical for the success of the upgrade. Section 4 is a brief conclusion to the paper.

2 Context and Background

Many large organizations involved in various critical government roles rely on a Common Operations System (COS) for planning their business operations; there are some two dozen instances of the COS in operation around the country. The users of the COS depend on it for a large percentage of their business tasks. The responsibility for developing and sustaining the COS for all communities of users is given to a single government agency, called herein the COS Development Agency (CDA), which is also one of the COS users.

By virtue of the manner in which it has evolved over the last two decades, the COS is a large aggregation of about 50 independently developed systems in each COS. Many of these constituent systems are themselves composed of other systems. The COS thus exemplifies the need, described in the previous chapter, to define unambiguous perspectives when discussing or analyzing systems of systems. (See Figure 1.) Further, while each COS instance has a basic set of applications, the instances are not identical. In each COS in operation, a varying number of “unofficial” applications, developed without the knowledge of the CDA, are unique to that particular COS.

The CDA is presently considering several courses of action aimed at modernizing the COS. Of particular significance for this case study is a planned replacement of the COS’s key infrastructure system. There are other planned modernization activities that may also occur, including a program to migrate all of the COSs toward a single standardized baseline. To the extent that these other modernization activities are relevant to our study, they shall be described below. However, the principal topic of this paper will be the replacement of the key infrastructure portion of the COS.

2.1 Description of the COS

The COS consists of a large number of applications that all rely on a complex infrastructure. Examples of these applications include planning, distribution control, and rapid resupply. These applications execute such tasks as managing delivery operations, estimating outcomes of planned operations, and assisting in various security tasks. In carrying out its mission, the COS provides its users access to various types of data (e.g., the availability of certain assets), and performs risk/benefit analysis computations. In Figure 2, we depict the high-level shape (considerably simplified) of a COS; this takes the perspective of an overall COS as the system of systems, and its primary applications and infrastructure as the constituent systems. Note that we are only concerned with the interoperability relationships between each application and the infrastructure, not among the individual applications. The latter will be considered in greater detail in Section 3.

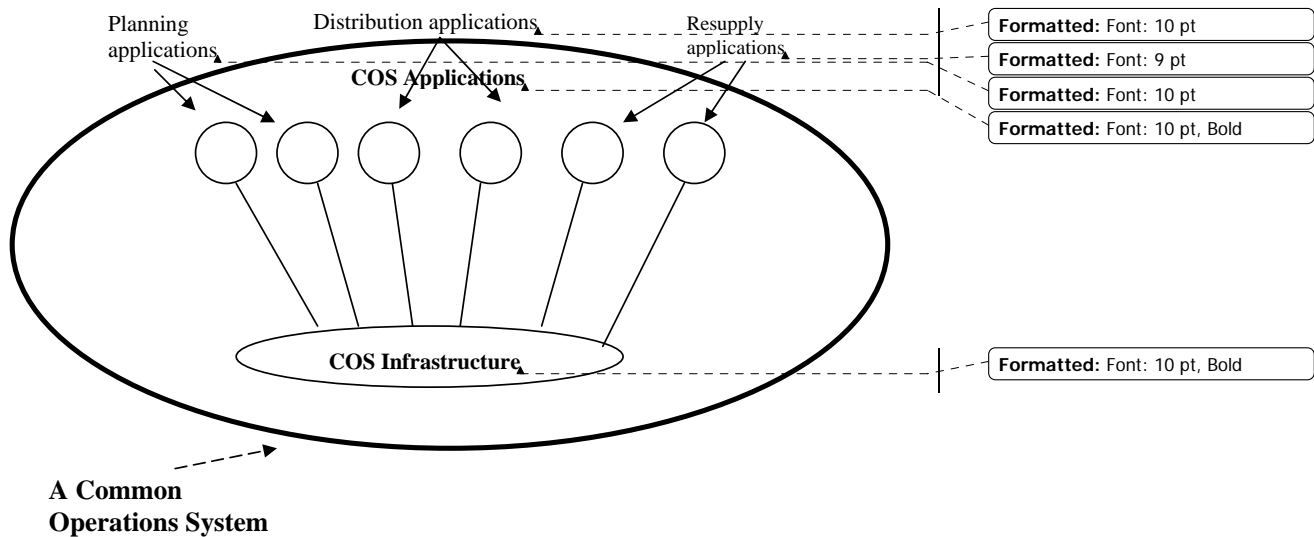


Figure 2: High-Level COS Organization

The most significant tasks performed by a COS are planning activities, particularly the generation of Distribution Orders (DOs). The DO is the primary mechanism for conducting operations for most users of the COS, and is a capability of primary importance. The generation of the DO by the present system requires a two-day period. The generation activities begin with a coordination meeting, move through a number of refinement activities and logistics plans, and end with execution roughly 48 hours later.

All the COS applications, and the capabilities that they provide, depend in one manner or another on the COS's infrastructure, the Common Function system (CF), developed by Allied Industries, and fielded in 1995. The CF has been called the "engine" of the COS by many of its users, and the quality of its upgrade essentially means the quality of the COS for the future.²

At the heart of the CF system are four autonomous databases (called herein DB_a, DB_b, DB_c, and DB_d) that contain many data elements in common. Each of these databases is accessed by the various applications within the COS. Some of these applications use a service-based interface to these databases, but the majority of the database accesses are accomplished through SQL calls from the various applications. Using these four databases, and integrating data from the various applications, the CF provides information and decision support to combined planning, delivery, and oversight operations.

² The actual situation is even more complex, since the CF also provides some application services as well as its infrastructure services. However, this complexity does not change the essence of our analysis of the interoperability issues found in the study.

It is only the CF for which an upgrade is now being planned. None of the other elements in the COS is expected to be replaced in conjunction with the planned upgrade of the CF.

2.2 Organizations and Their Missions

Of the many organizations that are stakeholders in the COS and its modernization, the following are central.

- The COS Development Agency (CDA) is the large government organization with overall authority for developing, fielding, and maintaining the COS and its infrastructure.
- Within the CDA, the Integration and Fielding Group (IFG) has the specific responsibility for maintaining and fielding the COS for the entire COS community (i.e., both for the CDA and all the other COS users).
- Also within the CDA, the Rapid Experiment Group is the primary sponsor of the infrastructure upgrade and also speaks for the COS's users in the field.
- Allied Industries (AI), the original prime contractor for the existing CF, is still the prime contractor for the overall integration of the COS.
- Innovative Data Systems is the developer of a widely used data fusion tool (DFuse). This tool was modified to become the planned upgrade for the CF.

2.3 Interoperability in the COS

In Section 1, we defined interoperability as

the ability of a collection of communicating entities to (a) share specified information and (b) operate on that information according to a shared operational semantics in order to achieve a specified purpose in a given context.

This definition is clearly applicable to the COS: the system's purpose is to unite data from a large number of independent, stovepiped systems and thereby provide its users with a unified means for making and executing plans for complex operations.

To accomplish the data sharing, and hence to successfully support the COS's mission, the services of the CF are the principal mechanisms that create the necessary interoperability. These integrating capabilities include message parsing, common GUI services, a unified protocol for data storage and retrieval, and comparable tasks. The operations of the various applications are otherwise independent; their integration is achieved only through calls to the CF. Thus, the CF provides the interoperability relationships we defined in Section 1; they are the "lines between the circles" seen in Figure 1.

The COS's interface issues are actually more varied than this description suggests, since there are also many other systems external to the COS with which the COS must interoperate. The interfaces to these systems are primarily message oriented and are carried out independently

by many of the individual COS applications as well as by the CF. The message formats used are of several standardized formats and are received by one or more applications. The messages may simply supply data to the COS or may require processing and response. These interactions with other non-COS systems make the interoperability issues of the COS especially complex and make the upgrade of its infrastructure all the more difficult a task.

2.4 Description of the Planned Upgrade

The upgrade of the CF has been under consideration for some years, due to a growing perception that the present state of the CF is insufficient, for a number of reasons. First, the CF contains multiple databases, with overlapping and redundant functionality. For instance, much of the data in DB_a is duplicated by data in the other three databases. Second, no uniform user interface exists among the applications; the various systems appear fragmented and piecemeal to the user. Third, the CF currently does operations planning using a phased planning-execution model, which is usually employed in a three-stage process; it is not capable of performing dynamic planning and replanning, nor planning of multiple simultaneous operations. Finally, within the past few years, the entire government has made a strong effort to migrate all of its systems to reflect modern computational practice. For major systems like the COS, it is now virtually a mandate that they should be Web enabled, which the current system is not. There is, therefore, growing agreement that some modernization of the CF is needed. But the precise details of that modernization, and who should execute it, have only recently been defined.

The CDA funds a number of experimental initiatives through its Rapid Experimental Group (REG). One of these initiatives, in existence for three years, has been an experiment to determine how best to modernize the CF. This experimental program has been carried out by Intelligent Data Systems (IDS), a small, inventive software company. One reason for this choice is that IDS has, over roughly 10 years, developed and matured a product called DFuse, a tool that reads data from multiple sources and presents it in a single unified view for the user. DFuse is currently in use in several large industry consortiums ranging from pharmaceutical manufacture to online auction support. The capability of DFuse is very similar to the desired capabilities of a modernized CF and some of its applications.

The experimental initiatives sponsored by REG are usually tested through a hands-on, try-it-out approach, one that is quite different from traditional developmental test/operational test (DT/OT). Users' impressions and reactions are paramount, and painstaking comparison with a large set of requirements is less critical. (As is logical, since the experiments are typically not following a predefined set of requirements in any case.) The experimental testing approach generally involves a considerably shorter test cycle than traditional government testing.

Within the past six months, a proposal was made by REG to convert the CF modernization initiative to a formal development; the work of IDS would now result in the actual CF upgrade. This proposal was accepted. However, it was also decided, in order to give end users

access to the improvement as quickly as possible, that the planned testing of the CF upgrade would not be done by a traditional DT/OT cycle but instead through the more informal testing approach that had originally been planned. To reduce the risks of this strategy, REG proposed modifying the experimental testing approach somewhat, to incorporate some DT and OT procedures. However, the compressed time of testing would remain.

Testing of the CF upgrade is expected to occur in early 2006, and a new version of the COS, with the upgraded CF, is to be fielded in mid-2006.

3 Specific Interoperability Issues

In this section we examine in detail many of the interoperability issues surrounding the planned replacement of the CF. We consider these issues in terms of risks to the interoperability between the CF and the COS applications, as well as risks relating to the interoperability between the COS and other external systems.

Note that the interoperability risks described below are of different types. Some are more technical in nature (i.e., machine-machine); others are more organizational (i.e., human-human); and still others are some combination of both.

3.1 System Interfaces

The CF has interfaces to most of the applications within the COS. In addition, the CF is visible to and has many interfaces to systems external to the COS. The majority of these interfaces involve the largest database (DB_a), which is used to create the DOs and is critical to many other COS operations. Some interfaces, particularly those with external systems, are in the form of messages whose format is either standardized or is specified through an ad hoc agreement. Other interfaces, largely those between the CF and the internal COS applications, make use of either SQL or service-based interfaces.

In upgrading the CF, IDS has completely redesigned DB_a, has eliminated the other three databases, has subsumed the functionality of some applications, and has introduced a wholly new set of services as the sole interface; the SQL interface will no longer be available. The developers have promised **exact** support for a small number of the existing services; they also promise **close** (but not exact) support for some others. They also warn that there will be **no** support for the remainder of the existing interfaces.

3.1.1 Elements of Risk

There are two principal risks in this area, and both of them pertain to machine-machine interoperability. The first is that the developer of the CF replacement has made unilateral changes to its interfaces. The second is the existence of the “unofficial” applications in the different COSs throughout the country.

Unilateral Changes to Interfaces: In a large-scale system of systems, a change in interfaces typically involves agreement among all organizations responsible for the systems that provide and use the interface. However, in the case of the CF upgrade, it appears that IDS has unilaterally defined a new set of services. Further, they have decreed that these services will

now be the only way to interface with the CF (i.e., SQL interfaces will no longer work). While this strategy may be technically appropriate, it brings with it a number of unknowns:

- It is unknown, in many cases, who has the responsibility for modifying the external systems to meet the new interfaces. This task is especially complex given the spectrum from COS internal applications to systems completely outside the COS community.
- It is unknown how much work any of these organizations must perform to make the necessary changes.
- The schedule for these other organizations to make their changes is unknown.

Unofficial Applications: Another area of risk is that, even if IDS were able to consult with all organizations responsible for systems that are **known** to interface with the CF, there are other, unofficial systems that interface with the CF. It is understood by all sides of the COS community that the various organizations using the COS have developed small-scale programs of their own; many of these are quite important in assisting users in performing their tasks. The extent to which those programs rely on knowledge of the CF internals (e.g., the structure of DB_a) will greatly affect the ability to migrate those applications to the CF replacement. At present, there is no knowledge of how great this problem will be.

3.1.2 Observed Mitigations

There are some activities underway that mitigate the first risk (i.e., the unilateral changes to the CF interfaces). For one thing, the COS integration contractor, AI, is working to create and maintain an integrated schedule for the appropriate upgrade of all systems that interface with the CF. This schedule will be useful in many ways. For instance, it will facilitate the development of a pragmatic test plan (e.g., some interfaces may be testable without needing **all** interfaces to be tested) that includes full operational testing. Furthermore, such an integrated schedule makes explicit the commitments that the various other organizations are making to the upgrade. Should the schedule indicate that requisite changes will not be made on time, then the CF or COS management negotiations might be undertaken to facilitate some means to accelerate the other organizations' schedule.

Another mitigating activity concerns efforts to “socialize” members of the COS community. When a change is made to the interfaces as dramatic as that in the CF upgrade, significant resistance is likely, and from many quarters. Thus, dissemination of information about the expected level of improvement is critical. To that end, IDS has been holding events at which the improved functionality and “look and feel” of the CF upgrade are demonstrated, in the expectation that users who attend can not only provide feedback to the developers, but also act as champions for the upgrade when they return to their home organizations.

This mitigation does not fully address the risk, however. For those organizations that develop and maintain COS applications, the socialization efforts are a useful mechanism to acquaint their personnel with information about the new interfaces. But for those external systems that

simply communicate with the COS (or with the CF directly), the probability that such information dissemination will lead to the necessary modifications is sharply reduced.

In any case, the socialization activity is also a mitigation for the second risk noted above (i.e., the existence of unofficial applications), since it is likely that those converted to a favorable opinion of the upgrade will also be proactive in describing the advantages of the upgrade back in their home organizations. These affected organizations will then be more likely to prepare for the upgrade, including making plans for how the unofficial applications are to be upgraded to accommodate the changed CF interfaces.

3.2 Organizational Responsibilities

For the CF upgrade to succeed, it is necessary that the four organizations involved in the development communicate and cooperate. Currently, the organizational structure is such that the REG has authority over IDS, and IFG has authority over AI. REG and IFG are both subgroups within the CDA, but in the large, complex CDA organization, it is not well defined where the respective authority of REG and IFG starts and ends. By contrast, the two industrial organizations, AI and IDS, have a contractual relationship for cooperative work and also an informal working agreement. Figure 3 shows these relationships; the solid arrows represent relationships of authority and reporting and the dashed arrows represent relationships of cooperation.




Figure 3: Organizational Links

As noted earlier, REG represents the CF user community and has been the primary sponsor of the experimental program on possible modernizations to the CF. IDS has been the software company that conducted those experiments, resulting in an experimental CF upgrade, now chosen to become the actual CF upgrade. IFG has responsibility for delivering the COS (and hence the CF) to the user community. The AI corporation developed the existing CF and has been (and will continue to be) the overall integrator of the COS.

3.2.1 Elements of Risk

There are two principal risks in this area, and both of them pertain to organizational interrelationships. The first is that there appears to be a clash of goals between the two suborganizations of the CDA, REG and IFG. The second is that the contractual relationship between the industry corporations, IDS and AI, is “prime-prime.”

Clash of Goals: REF and IFG have different views of what is best for the COS community. REG regards its mission as the dissemination of new capabilities throughout the community as quickly as possible, to maintain agility in the current climate of rapid information technology development. REG personnel therefore value speed and inventiveness and are not averse to taking on significant risk to improve the condition of IT-related capabilities for end users. IFG, on the other hand, has legal responsibility for preserving the overall health and stability of the major CDA systems (which includes monitoring the health and stability of the companies that build its components). It also deals with the other government agencies that use the COS system. IFG personnel value procedure, validation, and comparable virtues. This makes for tension between the two: REG would like to cause new functionality to be developed and fielded immediately; IFG prefers more careful development, with full testing and assurance that the capability can be sustained, with reasonable cost, over the lifetime of its fielding. This mode of development is necessarily slower than REG would prefer.

Prime-Prime Contractual Relationship: A second area of risk arises from the nature of the prime-prime relationship between IDS and AI. Under this relationship, IDS is the prime contractor for creating the upgraded version of the CF, and AI is the prime contractor with responsibility for integrating the upgraded CF into the COS (i.e., with the other infrastructure components and the applications). This poses the risk that responsibility for some tasks may be undefined, introducing the occasion for “finger pointing.” For instance, AI may at some point determine the need for some change in the CF. Since IDS is the prime contractor for **developing** the CF upgrade, AI has no authority to demand that any changes be made (as would occur if AI were a prime and IDS a sub). Under the prime-prime relationship, AI can only describe the need for the change to IDS and hope that IDS will implement the change as requested.

Note that this risk is only exacerbated by the lines of authority depicted in Figure 3. If AI sees the need to demand the hypothetical change, AI must request it from IFG; IFG must negotiate it with REG, who then would instruct IDS to make the desired change. This type of organizational structure has many drawbacks, not least of which is that it virtually guarantees that the expected schedule will not be met.

3.2.2 Observed Mitigations

We saw few mitigations in place regarding the clash of goals between REG and IFG. Because of their unresolved differences, many activities that should have taken place by now (e.g., finalizing a contract for AI with IDS as a subcontractor for integration and support) have not

yet been performed as of this writing. These delays will, in the long run, delay the program, which is already on an aggressive schedule.

Regarding IDS and AI, the relationship between the two is better than the relationship between REG and IFG. The two companies appear to respect each other's abilities and have, to some extent, arrived at an appropriate division of effort (IDS acting as the software development organization and AI acting as the system engineering organization). Further evidence of this relationship is the team formed by the two companies to work on another (i.e., non-CF) project.

3.3 Requirements and Functionality

The system that will become the upgraded CF was originally developed as an experiment to determine whether some of the drawbacks of the current system might be improved. From its original conception even to the point of our observations of the program, the experimental system had no specific requirements, other than to see what improvements might be feasible. The decision to convert that experimental system into a fieldable system of record was made relatively recently. That mode of development has resulted in some anomalous circumstances.

For instance, since IDS began its work as an experiment, it did not work from a set of requirements, but only investigated certain likely directions of desirable improvement. Although its task has been converted from an experiment to developing a system of record, there were still, as of our study, no requirements that IDS works to. This is a source of concern, because AI, which developed the original CF (and which will integrate the new version into the COS), did so with a well-defined set of requirements, both high level and low level. These requirements are still applicable to the existing CF, and so far as we can determine, it is expected, but not clearly defined, that these present requirements will also apply to the CF upgrade.³

In addition to any formal requirements, there is also commitment from all parties to "delivering the same (or equivalent) capabilities" in the upgrade. However, there are some anomalies here as well. For instance, it was mentioned by all parties that the new system's users will be able to perform all the functions of the existing system, but "not necessarily in the same way." Given that the upgrade will offer significant changes (e.g., to the user interface) and will streamline at least some of the current redundancies, it is not clear whether this commitment is fully understood in the same manner by IDS, AI, REG, and IFG.

³ IDS has just informally agreed to take responsibility for many of the high-level requirements, but the precise manner in which these requirements will be met is not yet clear.

3.3.1 Elements of Risk

The most significant area of risk to the program is that, as of this writing, IDS has not yet formally signed up to any of the requirements for the CF. Given the contractual relationship between AI and IDS, mitigating this risk will require that their respective authorities (REG and IFG) mutually assure that all existing CF requirements will be met by the upgraded CF. For this to occur, it is necessary that communications that pass through from AI to IFG to REG to IDS (and the reverse) are interpreted consistently.

A second area of risk is that the ultimate user community (and not just their REG representatives) may reject the CF upgrade because it “doesn’t work the way we’re used to.” This risk is amplified because, given the accretive development of the COS, there are many undocumented features in the CF.

3.3.2 Observed Mitigations

IDS is doing a good job of involving as many users as possible in early demonstrations of the CF upgrade. IDS listens to user feedback and incorporates changes into the system as much as possible, thus reducing the risk of user rejection. However, beyond this we saw few mitigations to the issue of requirements being unmet (at least in the eyes of some users) by the CF upgrade.

3.4 Development and Integration Processes

IDS and AI each employs a robust set of development and integration processes. From each company’s individual standpoint, its respective processes are well suited to its own product domains, and to its customers’ expectations. IDS uses relatively lightweight development processes, with emphasis on hands-on, iterative development to quickly provide desired functionality driven by a database of internally agreed-upon issues. AI uses its extensive corporate processes to guide very deliberate system engineering and development practices. Further, IDS has more of an application focus, while AI takes more of a systems view.

3.4.1 Elements of Risk

The risk in this area pertains to organizational interoperability. It stems from the difference in the software development and system integration processes of IDS and AI. Although the risk is described in terms of process, it is manifested in other ways: in corporate culture and even in development tools. For example, the two companies use different tools for tracking work and managing requirements; this complicates the process of exchanging information.

It is not known if past collaborations (some years ago, the two companies collaborated in a previous integration effort) are indicators of behavior in this case because the differences between that effort and the present integration effort are quite significant. The CF upgrade integration will be far more complex, with close coupling between the new infrastructure and

all of the other components in the COS. The previous effort simply involved taking the IDS component as a “black box” with a clearly defined interface to (the existing) CF.

3.4.2 Observed Mitigations

One successful strategy is mitigating the potential risks arising from the process mismatch: both companies are working closely together. Each takes proactive steps to understand what the other is doing, why it is doing it, and how that relates to its own process. They have thus already taken steps to ensure visibility into each other’s processes.

Another mitigating factor is AI’s confidence in IDS software. Good software is easier to integrate than poorly written software, and AI’s previous experience with the integration of an IDS product into the CF created trust in the quality of software developed by IDS and the processes used to manage that development.

3.5 Testing

There are several innovative aspects of the COS infrastructure modernization effort. One of these results from the experimental environment in which the CF upgrade was developed. It is common, in such experimental efforts, to perform testing in an equally experimental manner. Users are encouraged to try out the system and examine its features; there is not commonly a set of requirements to which the system is tested.

At least some of the testing of the CF upgrade will be done in this manner. There will be some more traditional testing, but it is not entirely clear how much of the testing will be rigorous OT/DT and how much will be on a more informal, “let’s try it out” basis. While this approach has the potential to reduce the time to field the COS infrastructure modernization, there are several potential risks, all of which are interrelated, and all of which include both technical and organizational aspects.

3.5.1 Elements of Risk

Division of responsibility: The division of responsibility noted above in Section 3.2.1 (e.g., the prime-prime relationship) poses risks in the testing area as well. As the CF upgrade is being developed and incrementally released to AI, test problem reports (TPR) are generated by AI. If a TPR has a resolution that is ambiguous, the same kind of risks noted above can arise. If this should occur later, during the experimental testing period, it may result in delays in a testing schedule that is already condensed.

Immature test procedures: IDS’s test procedures are relatively immature: every activity is at or below Level A, the lowest defined level in the Test Process Improvement (TPI) scale. There is no automated testing in place yet, nor is there a capability to test the system under its expected operational load. The Quality Assurance Department is fairly new (less than one year old), and is minimally staffed.

Lack of clear requirements: We have already noted that the CF upgrade was developed without formal requirements. The experimental testing approach described above commonly uses “operational threads” to guide the participants, and the requirements for the existing CF are the most likely source for developing these to use for testing the CF upgrade. But there is considerable uncertainty over the relationship between these operational threads and any requirements that apply to the CF upgrade.

Combining testing approaches: Performing traditional testing in conjunction with an experimental testing approach is attractive from the perspective of saving time. But the demands of a successful DT and OT, in the traditional approach, are quite different. In an experimental approach, considerable time is spent just getting the participating systems to work together; detailed isolation of problems is challenging. Further, there is very little time in an experimental testing approach to correct any problems, conduct regression testing, and integrate new baselines.

3.5.2 Observed Mitigations

The division of responsibility between AI and IDS is being addressed—in part—by close cooperation between AI and IDS. Whereas AI is responsible for the initial screening on TPRs, IDS supports that activity, thus lessening the probability of an issue “falling between the cracks.”

One factor that mitigates some testing risk is that software produced by IDS is apparently of high quality. This diminishes the need to depend on Quality Assurance to “test quality in.” However, we were still dismayed to observe little evidence that the relative immaturity of the IDS test organization and test procedures is being addressed; doing so would include making changes to the corporate culture within IDS.

The lack of clear requirements, coupled with the combining of the experimental testing approach and formal testing, are areas of very high risk with little apparent mitigation in place. We are particularly concerned that in the absence of formal requirements the testers for OT will independently define their view of the function of the CF, increasing the likelihood that the CF will fail OT.

4 Summary

This technical note has described how the infrastructure upgrade poses a number of risks to the interoperation of the COS. These include risks in the purely technical domain (i.e., machine-machine) as well as in the organizational domain (i.e., human-human). The former include unilateral changes to the infrastructure's interfaces and the lack of definition of both requirements and functionality. The latter include conflicting goals of the sponsoring organizations and incompatible processes in the engineering organizations.

However, the major risk lies in the area of testing. From the point of view of particular risks in the planned testing strategy, the division of responsibility discussed in Section 3.2.1 is a major concern, as is the lack of clear requirements to govern the testing. But from a broader point of view, the strategy to test so complex a system, and one on which so many other systems depend, in so informal a manner as currently planned, poses a grave risk to the continued successful operation of the COS.

References

URLs are valid as of the publication date of this document.

- [Carney 05]** Carney, D.; Fisher, D.; Morris, E.; & Place, P. *Some Current Approaches to Interoperability* (CMU/SEI-2005-TN-033). Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University, 2005. <http://www.sei.cmu.edu/publications/documents/05.reports/05tn033.html>.
- [Maier 98]** Maier, Mark. *Architecting Principles for Systems of Systems*. <http://www.infoed.com/Open/PAPERS/systems.htm> (1998).

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE November 2005		3. REPORT TYPE AND DATES COVERED Final
4. TITLE AND SUBTITLE Topics in Interoperability: Infrastructure Replacement in a System of Systems			5. FUNDING NUMBERS FA8721-05-C-0003	
6. AUTHOR(S) David Carney, James Smith, Patrick Place				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Software Engineering Institute Carnegie Mellon University Pittsburgh, PA 15213			8. PERFORMING ORGANIZATION REPORT NUMBER CMU/SEI-2005-TN-031	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) HQ ESC/XPK 5 Eglin Street Hanscom AFB, MA 01731-2116			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12A DISTRIBUTION/AVAILABILITY STATEMENT Unclassified/Unlimited, DTIC, NTIS			12B DISTRIBUTION CODE	
13. ABSTRACT (MAXIMUM 200 WORDS) This technical note examines the Common Operations System (COS), a large aggregation of independently developed systems, and the risks posed to it by an infrastructure upgrade. Many large organizations involved in various critical government roles depend on the COS for planning their business operations. When such a large number of applications rely on a complex infrastructure, an attempt to upgrade raises many interoperability issues. The risks involved, and their observed mitigations, are examined in several areas: system interfaces, organizational responsibilities, requirements and functionality, developing an integration process, and testing.				
14. SUBJECT TERMS interoperability, system of systems, SoS, infrastructure upgrade			15. NUMBER OF PAGES 29	
16. PRICE CODE				
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	